

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:

TI-35452

Gerard Chauvel, et al.

Art Unit: 2183

Serial No: 10/632,222

Examiner:

Filed: July 31, 2003

Conf. No.: 2081

For: Program Counter Adjustment Based on the Detection of an Instruction Prefix

**TRANSMITTAL LETTER ACCOMPANYING CERTIFIED COPY OF
PRIORITY APPLICATION UNDER 35 U.S.C. §119**


Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

MAILING CERTIFICATE UNDER 37 C.F.R. §1.8(a)

I hereby certify that the above correspondence is being deposited with the U.S. Postal Service as First Class Mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on March 12, 2004.


Robin E. Barnum

Dear Sir:

Submitted herewith is a certified copy of **European** Patent Application No. **03291920.1** (**TI-35452EP**), filed on **July 30, 2003**, in the **European** Patent Office and from which priority under 35 U.S.C. §119 is claimed for the above-identified application.

Respectfully submitted,



Robert D. Marshall, Jr.
Attorney for Applicant
Reg. No. 28,527

Texas Instruments Incorporated
P.O. Box 655474, MS 3999
Dallas, TX 75265
(972) 917-5290



100-100000-100000



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

TL-35452EP
FOR: S/N 10/632,222
(TL-35452)

Bescheinigung

Certificate

Attestation

Die angehefteten Unterla-
gen stimmen mit der
ursprünglich eingereichten
Fassung der auf dem näch-
sten Blatt bezeichneten
europäischen Patentanmel-
dung überein.

The attached documents
are exact copies of the
European patent application
described on the following
page, as originally filed.

Les documents fixés à
cette attestation sont
conformes à la version
initialement déposée de
la demande de brevet
européen spécifiée à la
page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

03291920.1

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk



Anmeldung Nr:
Application no.: 03291920.1
Demande no:

Anmeldetag:
Date of filing: 30.07.03
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

Texas Instruments Incorporated
7839 Churchill Way
Dallas,
Texas 75251
ETATS-UNIS D'AMERIQUE
TEXAS INSTRUMENTS FRANCE
Avenue Bel Air
06271 Villeneuve Loubet Cédex
FRANCE

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se référer à la description.)

Program counter adjustment based on the detection of an instruction prefix

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s)
revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

US/31.07.02/US 400391 P

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

G06F9/00

Am Anmeldetag benannte Vertragsstaaten/Contracting states designated at date of
filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IT LU MC NL
PT RO SE SI SK TR LI

PROGRAM COUNTER ADJUSTMENT BASED ON THE DETECTION OF AN
INSTRUCTION PREFIX

The present invention relates generally to processors and more particularly to a processor capable of changing a decoder behavior based on the detection of a prefix during a pre-decode step.

Many types of electronic devices are battery operated and thus preferably consume as little power as possible. An example is a cellular telephone. Further, it may be desirable to implement various types of multimedia functionality in an electronic device such as a cell phone. Examples of multimedia functionality may include, without limitation, games, audio decoders, digital cameras, etc. It is thus desirable to implement such functionality in an electronic device in a way that, all else being equal, is fast, consumes as little power as possible and requires as little memory as possible. Improvements in this area are desirable.

BRIEF SUMMARY OF THE PREFERRED EMBODIMENTS

As disclosed herein, a processor (e.g., a co-processor) includes a decode logic decoding a current instruction in parallel with a pre-decode logic pre-decoding the next instruction to determine if the next instruction includes a predetermined prefix. The prefix may indicate an instruction set to which the next instruction belongs. The prefix also may indicate the format of the next instruction. In accordance with at least some embodiments of the invention, the decode logic may change the decoding behavior of the decode logic based on the type of prefix determined by the pre-decode logic. In addition, the pre-decode logic also may further determine the identity of subsequent bytes in parallel with the decoding of the current instruction by the

decode logic. In some embodiments, the pre-decode logic may pre-decode at least subsequent five bytes in parallel with the decoding of the current instruction.

Other embodiments may include a system (e.g., a cellular telephone) that includes a main processor unit coupled to a co-processor. The co-processor may include a decode logic. The decode logic may be able to decode instructions from different instruction set and may also decode different instruction formats (e.g., variable instruction lengths or variable opcode lengths). In addition, the co-processor may also include a pre-decode logic coupled to the decode logic. The pre-decode logic operates in parallel with the decode logic, in which the pre-decode logic pre-decodes a next instruction while the decode logic decodes a current instruction. The pre-decode logic determines what instruction set the next instruction belongs to and/or the format of the next instruction. In particular, the pre-decode logic determines if the next instruction consists of a prefix. If a prefix is determined, the decode logic may temporarily change the behavior of the decoding operation for the next instruction.

In some embodiments, a register (e.g., a program counter) may be used to indicate a current instruction the decode logic may be decoding. However, in cases where a prefix of a next instruction is determined by the pre-decode logic, the program counter may skip the prefix, such that the decode logic does not decode the prefix. As such, the prefix does not enter the normal execution pipeline flow and therefore, no time penalty is incurred for executing two instruction sets or multiple instruction formats.

NOTATION AND NOMENCLATURE

Certain terms are used throughout the following description and claims to refer to particular system components. As one skilled in the art will appreciate, computer companies may refer to a component by different names. This document does not intend to distinguish between components that differ in name but not function. In the following discussion and in the claims, the terms "including" and "comprising" are used in an open-ended fashion, and thus should be interpreted to mean "including, but not limited to...". Also, the term "couple" or "couples" is intended to mean either an indirect or direct connection. Thus, if a first device couples to a second device, that connection may be through a direct connection, or through an indirect connection via other devices and connections.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more detailed description of the preferred embodiments of the present invention, reference will now be made to the accompanying drawings, wherein:

Figure 1 shows a diagram of a system in accordance with preferred embodiments of the invention and including a Java Stack Machine ("JSM") and a Main Processor Unit ("MPU");

Figure 2 shows a block diagram of the JSM of Figure 1 in accordance with preferred embodiments of the invention;

Figure 3 shows various registers used in the JSM of Figures 1 and 2;

Figure 4 illustrates decoding and pre-decoding instructions in parallel; and

Figure 5 shows a block diagram of a pre-decode logic usable in the JSM.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The following discussion is directed to various embodiments of the invention. Although one or more of these embodiments may be preferred, the embodiments disclosed should not be interpreted, or otherwise used, as limiting the scope of the disclosure, including the claims, unless otherwise specified. In addition, one skilled in the art will understand that the following description has broad application, and the discussion of any embodiment is meant only to be exemplary of that embodiment, and not intended to intimate that the scope of the disclosure, including the claims, is limited to that embodiment.

The subject matter disclosed herein is directed to a programmable electronic device such as a processor. The processor described herein may be particularly suited for executing Java™ Bytecodes, or comparable, code. As is well known, Java is particularly suited for embedded applications. Java is a relatively "dense" language meaning that on average, each instruction may perform a large number of functions compared to many other programming languages. The dense nature of Java is of particular benefit for portable, battery-operated devices that preferably include as little memory as possible to save space and power. The reason, however, for executing Java code is not material to this disclosure or the claims that follow. The embodiment of the invention may be described in the context of Java but should not limit the execution of only Java instructions. The processor described herein may be used in a wide variety of electronic systems.

Referring now to Figure 1, a system 100 is shown in accordance with a preferred embodiment of the invention. As shown, the system includes at least two processors 102 and 104. Processor 102 is referred to for purposes of this disclosure as a Java

Stack Machine ("JSM") and processor 104 may be referred to as a Main Processor Unit ("MPU"). System 100 may also include memory 106 coupled to both the JSM 102 and MPU 104 and thus accessible by both processors. At least a portion of the memory 106 may be shared by both processors meaning that both processors may access the same shared memory locations. Further, if desired, a portion of the memory 106 may be designated as private to one processor or the other. System 100 also includes a Java Virtual Machine ("JVM") 108, compiler 110, and a display 114. The JSM 102 preferably includes an interface to one or more input/output ("I/O") devices such as a keypad to permit a user to control various aspects of the system 100. In addition, data streams may be received from the I/O space into the JSM 102 to be processed by the JSM 102. Other components (not specifically shown) may be included as well.

As is generally well known, Java code comprises a plurality of "Bytecodes" 112. Bytecodes 112 may be provided to the JVM 108, compiled by compiler 110 and provided to the JSM 102 and/or MPU 104 for execution therein. In accordance with a preferred embodiment of the invention, the JSM 102 may execute at least some, and generally most, of the Java Bytecodes. When appropriate, however, the JSM 102 may request the MPU 104 to execute one or more Java Bytecodes not executed or executable by the JSM 102. In addition to executing Java Bytecodes, the MPU 104 also may execute non-Java instructions. The MPU 104 also hosts an operating system ("O/S") (not specifically shown), which performs various functions including system memory management, the system task management that schedules the JVM 108 and most or all other native tasks running on the system, management of the display 114, receiving input from input devices, etc. Without limitation, Java code may be used to perform any one of a variety of applications including multimedia, games or web-based applications in the system 100, while non-Java code, which may

comprise the O/S and other native applications, may still run on the system on the MPU 104.

The JVM 108 generally comprises a combination of software and hardware. The software may include the compiler 110 and the hardware may include the JSM 102. The JVM may include a class loader, Bytecode verifier, garbage collector, and a Bytecode interpreter loop to interpret the Bytecodes that are not executed on the JSM processor 102.

In accordance with preferred embodiments of the invention, the JSM 102 may execute at least two instruction sets. One instruction set may comprise standard Java Bytecodes. As is well known, Java is a stack-based programming language in which instructions generally target a stack. For example, an integer add ("IADD") Java instruction pops two integers off the top of the stack, adds them together, and pushes the sum back on the stack. In general, the JSM 102 comprises a stack-based architecture with various features that accelerate the execution of stack-based Java code.

Another instruction set executed by the JSM 102 may include instructions other than standard Java instructions. In accordance with at least some embodiments of the invention, such other instruction set may include register-based and memory-based operations to be performed. This other instruction set generally complements the Java instruction set and, accordingly, may be referred to as a complementary instruction set architecture ("C-ISA"). By complementary, it is meant that the execution of more complex Java Bytecodes may be substituted by micro-sequences using C-ISA instructions that permit address calculation to readily "walk through" the JVM data structures. Further, such micro-sequences may also use Bytecode instructions. The execution of Java may be made more efficient and run faster by

replacing some sequences of Bytecodes by preferably shorter and more efficient sequences of C-ISA instructions. Bytecodes may also be used within a C-ISA sequence. The two sets of instructions may be used in a complementary fashion to obtain satisfactory code density and efficiency. As such, the JSM 102 generally comprises a stack-based architecture for efficient and accelerated execution of Java Bytecodes combined with a register-based architecture for executing register and memory based C-ISA instructions. Both architectures preferably are tightly combined and integrated through the C-ISA.

Figure 2 shows an exemplary block diagram of the JSM 102. As shown, the JSM includes a core 120 coupled to data storage 122 and instruction storage 130. The core may include one or more components as shown. Such components preferably include a plurality of registers 140, three address generation units ("AGUs") 142, 147, micro-translation lookaside buffers (micro-TLBs) 144, 156, a multi-entry micro-stack 146, an arithmetic logic unit ("ALU") 148, a multiplier 150, decode logic 152, instruction fetch logic 154, and pre-decode logic 158. In general, operands may be retrieved from data storage 122 or from the micro-stack 146, processed by the ALU 148, while instructions may be fetched from instruction storage 130 by fetch logic 154, pre-decoded by pre-decode logic 158, and decoded by decode logic 152. The address generation unit 142 may be used to calculate addresses based, at least in part on data contained in the registers 140. The AGUs 142 may calculate addresses for C-ISA instructions as will be described below. The AGUs 142 may support parallel data accesses for C-ISA instructions that perform array or other types of processing. AGU 147 couples to the micro-stack 146 and may manage overflow and underflow conditions in the micro-stack preferably in parallel. The micro-TLBs 144, 156 generally perform the function of a cache for the address translation and memory protection information bits that

are preferably under the control of the operating system running on the MPU 104. The decode logic 152 and the pre-decode logic 158 may be adapted to execute both the standard Java instructions as well as the C-ISA instructions of the system. In addition, the decode logic 152 and the pre-decode logic 158 may also be adapted to execute instructions of different format, including, but not limited to, variable opcode lengths and variable instruction lengths. The operation of the decode logic 152 and the pre-decode logic 158 will be described in more detail below.

Referring now to Figure 3, the registers 140 may include 16 registers designated as R0-R15. Registers R0-R3, R5, R8-R11 and R13-R14 may be used as general purposes ("GP") registers usable for any purpose by the programmer. Other registers, and some of the GP registers, may be used for specific functions. For example, registers R4 and R12 may be used to store two program counters. Register R4 preferably is used to store the program counter ("PC") and register R12 preferably is used to store a micro-program counter ("micro-PC"). In addition to use as a GP register, register R5 may be used to store the base address of a portion of memory in which Java local variables may be stored when used by the current Java method. The top of the micro-stack 146 is referenced in registers R6 and R7. The top of the micro-stack has a matching address in external memory pointed to by register R6. The values contained in the micro-stack are the latest updated values, while their corresponding values in external memory may or may not be up to date. Register R7 provides the data value stored at the top of the micro-stack. Registers R8 and R9 may also be used to hold the address index 0 ("AI0") and address index 1 ("AI1"). Register R14 may also be used to hold the indirect register index ("IRI"). Register R15 may be used for status and control of the JSM 102. As an example, one status/control bit (called the "Micro-Sequence-Active" bit) may indicate if the JSM 102 is executing a "simple"

instruction or a "complex" instruction through a "micro-sequence." This bit controls in particular, which program counter is used R4 (PC) or R12 (micro-PC) to fetch the next instruction. A "simple" Bytecode instruction is generally one in which the JSM 102 may perform an immediate operation either in a single cycle (e.g., an "iadd" instruction) or in several cycles (e.g., "dup2_x2"). A "complex" Bytecode instruction is one in which several memory accesses may be required to be made within the JVM data structures for various verifications (NULL pointer, array boundaries, etc.). Because these data structures are generally JVM-dependent and thus may change from one JVM implementation to another, the software flexibility of the micro-sequence provides a mechanism for various JVM optimizations now known or later developed.

Referring again to Figure 2, as noted above, the JSM 102 may be adapted to process and execute instructions from at least two instruction sets. One instruction set includes stack-based operations and the second instruction set includes register-based and memory-based operations. The stack-based instruction set may include Java Bytecodes. Java Bytecodes pop, unless empty, data from and push data onto the micro-stack 146. The micro-stack 146 preferably comprises the top n entries of a larger stack that may be implemented in data storage 122. Although the value of n may be vary in different embodiments, in accordance with at least some embodiments, the size n of the micro-stack may be the top eight entries in the larger, memory-based stack. The micro-stack 146 preferably comprises a plurality of gates in the core 120 of the JSM 102. By implementing the micro-stack 146 in gates (e.g., registers) in the core 120 of the processor 102, access to the data contained in the micro-stack 146 is generally very fast, although any particular access speed is not a limitation on this disclosure.

The second, register-based, memory-based instruction set may comprise the C-ISA instruction set introduced above. The C-ISA instruction set preferably is complementary to the Java Bytecode instruction set in that the C-ISA instructions may be used to accelerate or otherwise enhance the execution of Java Bytecodes.

The ALU 148 adds, subtracts, and shifts data. The multiplier 150 may be used to multiply two values together in one or more cycles. The instruction fetch logic 154 generally fetches instructions from instruction storage 130. The instructions may first be pre-decoded by the pre-decode logic by 158 and decoded by decode logic 152. Because the JSM 102 may be adapted to process instructions from at least two instruction sets, the decode logic 152 generally comprises at least two modes of operation, one mode for each instruction set. In particular, the decode logic unit 152 may include a Java mode in which Java instructions may be decoded and a C-ISA mode in which C-ISA instructions may be decoded. The decoding mode may be changed temporary to let a single instruction from the other instruction set to be executed in the mode set for the other instruction set. The instruction simply needs to be preceded by a byte such as the Java impdep1 byte. The pre-decode logic 158 determines if an instruction is preceded by this byte and consider this byte as an instruction prefix. The decode logic 152 and the pre-decode logic 158 may also decode instructions having variable formats based on the detection of a byte (e.g. wide) preceding these instructions, as will be discussed below.

The data storage 122 generally comprises data cache ("D-cache") 124 and data random access memory ("D-RAMset") 126. Reference may be made to copending applications U.S. Serial nos. 09/591,537 filed June 9, 2000 (atty docket TI-29884), 09/591,656 filed June 9, 2000 (atty docket TI-29960), and 09/932,794 filed August 17, 2001 (atty docket TI-31351). The stack (excluding the micro-stack

146), arrays and non-critical data may be stored in the D-cache 124, while Java local variables, critical data and non-Java variables (e.g., C, C++) may be stored in D-RAM 126. The instruction storage 130 may comprise instruction RAM ("I-RAM") 132 and instruction cache ("I-cache") 134. The I-RAMset 132 may be used for "complex" micro-sequenced Bytecodes or other "micro-sequences or critical sequences of codes," as will be described below. The I-cache 134 may be used to store other types of Java Bytecode and mixed Java/CISA instructions.

In accordance with the preferred embodiments, while the decode logic 152 decodes a current instruction, the pre-decode logic preferably decodes at least the first byte of the next instruction that follows the current instruction to determine whether the next instruction is preceded by predetermined prefixes. These prefixes are encoded on a byte in the preferred embodiment. The prefixes may be a single byte Java instruction (e.g., a Java wide instruction or a Java impdepl instruction) and may precede a next instruction, and thus, may be appended to the next instruction. Because in general the current instruction may be of a variable length, the first byte of the next instruction's Bytecodes may not be known a priori. Thus, the preferred implementation is to pre-decode enough Bytecodes ahead of the current value of the program counter to ensure that the beginning of the subsequent instruction is captured. In accordance with some embodiments, enough subsequent Bytecodes are pre-decoded to determine whether the next includes a predetermined prefix. In other embodiments, it is determined whether any of the instructions following the current instruction comprise a predetermined prefix, not just the next instruction.

Referring to Figure 4, a portion of a program may include a plurality of Bytes 162 (labeled as Bytes A-G). In Java, those bytes are called Bytecodes. Therefore, for simplicity, these

bytes may be referred as Bytecodes in both instruction sets. One or more of those bytes form a Java instruction. The Java instruction set, therefore, may include variable length instructions. In C-ISA, the instruction length may also vary in length (e.g., 2 to 4 bytes) and like the Java instructions, the length may only be determined during the decode. In accordance with the preferred embodiments, while the decode logic 152 is decoding Bytecode A, in which Bytecode A may be all or a portion of a current instruction, the next five Bytecodes preferably are provided to pre-decode logic 158. As illustrated, Bytecodes B through F are pre-decoded by the pre-decode logic 158 in parallel with the decoding of Bytecode A. The pre-decode logic 158 preferably determines if any of Bytecodes B through F include a predetermined "prefix." A prefix, as used herein, is a Bytecode that indicates the type of instruction that follows the prefix. For example, a format of an instruction, such as a "wide" instruction, may include a Java wide bytecode used as a prefix. In yet another example, a prefix (e.g., one of the Java implementation dependent and reserved code "Impdepl" Bytecode) may be used as a prefix and may indicate the presence of an instruction that belongs to a particular instruction set. As such, by first determining if an instruction differs in format or if the instruction belongs to a particular instruction set, the decode logic 158 may adjust the decoding behavior for decoding of the instruction.

In accordance with preferred embodiments of the invention, when the decode logic 152 is ready to begin decoding the next instruction, the program counter 160 (which may be stored in register R4, R12) may be incremented if the presence of a prefix in the next instruction was detected ahead of time by the pre-decode logic 158. In particular, the program counter may be incremented to skip the prefix and point to the next Bytecode in

the sequence 162. As such, time and processor resources are not expended decoding the predetermined prefix.

Referring again to the example of Figure 4, the decode logic 152 preferably decodes Bytecode A in parallel with the pre-code logic 158 pre-decoding Bytecodes B through F. During this pre-decode process, the pre-decode logic 158 may determine that Bytecode B may be preceded by a predetermined prefix (e.g., a Java "wide" prefix or a Java "impdep1" prefix). Once the decoding of Bytecode A is completed and Bytecode B has been determined by the pre-decode logic 158 to include the predetermined prefix, the program counter 160, currently pointing to Bytecode A, may be incremented to skip Bytecode B and point to Bytecode C. As such, when the decode logic 152 is finished decoding Bytecode A, the decode logic 152 decodes Bytecode C and skips Bytecode B.

In order to determine if a Bytecode comprises the predetermined prefix, the pre-decode logic 158 compares the m (e.g., five) Bytecodes following the Bytecode pointed to by the program counter with a predetermined set of bytes indicative of the predetermined prefix. Referring to Figure 5, the pre-decode logic 158 includes a plurality of comparators 200 and a multiplexer 202. In at least some embodiments, each of Bytecodes B through F may be compared to a predetermined prefix value. The prefix value may be equal to a Java wide prefix. A Java wide prefix, as described herein, may be an opcode that indicates the format of the instruction that immediately succeeds the prefix may be of a different format. More specifically, a Java wide prefix may extend the number of bits the succeeding instruction may contain. For example, a Java instruction, may include, but is not limited to, a parameter of eight bits. The Java wide prefix, may extend the Java instruction parameter to, for example, 16 bits. As such, to determine if the format of an instruction is different, each Bytecode B through F is compared to the Java wide prefix.

In some embodiments, a single Bytecode may match the Java wide prefix out of the five Bytecodes B through F. As such, if Bytecode C match the Java wide prefix, a high logic value, e.g., a logic "1", may be provided to multiplexer 202 via signal line W. A logic low value, e.g., a logic "0", may be provided to multiplexer 202 via signal lines V and X through Z. Based on the instruction length of the current instruction decoded by the decode logic 152, select line U may be provided to select from signal lines V through Z. For example, if Bytecode A is a single byte Java instruction, the decode logic provides a value n, where n is the number of Bytecode(s) that make up the current instruction decoded by the decoding logic 152. Since Bytecode A is a single byte Java instruction, select line U would select the next Bytecode, e.g., Bytecode B. If Bytecode B match the Java wide prefix, the program counter calculator 204, wherein the program counter calculator may be one or several adders, would increment the program counter 160 by $n + 1$. As such, the program counter 160 would skip Bytecode B and load the value of Bytecode C into the decode logic 152.

In some other embodiments, a plurality of the Bytecodes out of Bytecodes B through F may match the Java wide prefix. Similar to the process of determining the increment value of the program counter 160 for a single Bytecode match to the wide prefix, the program counter may be incremented based on the instruction length of the current instruction decoded by the decode logic 152. In particular, if Bytecodes B and F match the Java wide prefix, the program counter may be incremented to skip Bytecode B and subsequently skip Bytecode F.

Bytecodes B through F also may be compared to prefixes besides wide or to multiple prefix values. For example, Bytecodes B through F may be compared to a Java wide prefix and also to a

Java impdep prefix. A Java impdep prefix (e.g., impdep 1 or impdep 2) is an implementation dependent instruction. In a preferred embodiment, the Java impdep prefix may indicate that the succeeding instruction following the impdep prefix may be from a different instruction set as the previous instruction, such as a C-ISA instruction. Similar to the process for determining if Bytecodes B through F include a Java wide prefix, Bytecodes B through F also may be compared to a prefix value representative of a Java impdep prefix. Upon determining if at least one of the subsequent Bytecode is a Java impdep prefix, the program counter 160 may skip the prefix and load the value of the next Bytecode into the decoder. For example, while decoder 152 decodes Bytecode C of a current instruction, pre-decoder 158 may determine that Bytecode D is equal to the impdep prefix. Subsequently, the program counter 160 may skip over Bytecode D to Bytecode E. After the decoder 152 decodes Bytecode C, the decoder 152 decodes Bytecode E, without decoding Bytecode D.

In some embodiments, the sequential execution of the Bytecodes may be interrupted after a prefix has been pre-decoded by the pre-decode logic 158 but before the next instruction is loaded into the decode logic 152. For example, the decode logic 152 may currently decode a branch Bytecode which may interrupt the sequential execution of the plurality of bytes. Consequently, upon returning from the execution of the branch Bytecode, the prefix may be loaded into the decode logic 152. In this case, the prefix may be treated as a NOP because the prefix is presented as a first instruction without being pre-decoded, entering the decoder directly.

By predetermining if an instruction comprises a prefix, the behavior of decoder 152 may be modified to accommodate the instruction. Furthermore, by predetermining the prefix of a subsequent instruction in parallel with the decoding of a current

instruction, the decoder 152 may avoid decoding the prefix, improving the cycle time of decoding instructions of a different format and also instructions from a different instruction set.

While the preferred embodiments of the present invention have been shown and described, modifications thereof can be made by one skilled in the art without departing from the spirit and teachings of the invention. The embodiments described herein are exemplary only, and are not intended to be limiting. Many variations and modifications of the invention disclosed herein are possible and are within the scope of the invention. Accordingly, the scope of protection is not limited by the description set out above. Each and every claim is incorporated into the specification as an embodiment of the present invention.

CLAIMS

WHAT IS CLAIMED IS:

1. A processor, comprising:
 - instruction storage in which instructions are stored;
 - fetch logic coupled to the instruction storage to fetch instructions from the instruction storage;
 - decode logic coupled to the fetch logic to decode instructions fetched by the fetch logic; and
 - pre-decode logic associated with the decode logic;wherein at least some of the instructions comprise a prefix, and in parallel with the decode logic decoding a current instruction, the pre-decode logic determines whether a subsequent instruction comprise a prefix, in which case the decode logic causes a program counter to skip the prefix and changes behavior of the decode logic during decoding of the subsequent instruction.
2. The processor of claim 1, wherein at least some instructions comprise at least one Bytecode.
3. The processor of claim 1 or claim 2, wherein the prefix comprises a Java impdep instruction.
4. The processor of claim 3, wherein when detecting the Java impdep instruction, the subsequent instruction belongs to a different instruction set than the current instruction.
5. The processor of claim 1, wherein the prefix comprises a Java wide instruction.

6. The processor of any preceding claim, wherein when detecting the Java wide instruction changes format of the subsequent instruction.

7. The processor of any preceding claim, wherein in parallel with the decode logic decoding the current instruction, the pre-decode logic examines a predetermined number of subsequent bytes.

8. The processor of claim 7, wherein the predetermined number is at least 5.

9. A method of decoding variable length instructions, comprising:

decoding a current instruction according to a first behavior;

while decoding the current instruction, pre-decoding a subsequent instruction to determine if the subsequent instruction includes a predetermined prefix; and

if the subsequent instruction includes the predetermined prefix, causing a program counter to skip the predetermined prefix and changing the decoding of the subsequent instruction according to a second behavior.

10. The method of claim 9, wherein pre-decoding includes examining a predetermined number of bytes following the current instruction.

11. The method of claim 10, wherein the predetermined number is at least 5.

12. The method of claim 10 or claim 11, wherein pre-decoding further includes comparing each of the predetermined number of bytes to prefix value.
13. The method of claim 12, wherein the prefix value is equal to a Java impdep instruction.
14. The method of claim 12, wherein the prefix value is equal to a Java wide instruction.
15. The method of any of claims 9-14, wherein if the a Java wide prefix is detected, the first and second behaviors comprise a first mode for decoding instructions of a first format and a second mode for decoding instructions of a second format.
16. The method of any of claims 9-15, wherein if a Java impdep prefix is detected, the first and second behaviors comprise a first mode for decoding instructions of a first instruction set and a second mode for decoding instructions of a second instruction.

PROGRAM COUNTER ADJUSTMENT BASED ON THE DETECTION OF AN
INSTRUCTION PREFIX

ABSTRACT

A processor (e.g., a co-processor) comprising a decoder coupled to a pre-decoder, in which the decoder decodes a current instruction in parallel with the pre-decoder pre-decoding a subsequent instruction. In particular, the pre-decoder examines at least five Bytecodes in parallel with the decoder decoding a current instruction. The pre-decoder determines if a subsequent instruction contains a prefix. If a prefix is detected in at least one of the five Bytecodes, a program counter skips the prefix and changes the behavior of the decoder during the decoding of the subsequent instruction.

Figure 4




Figure 1

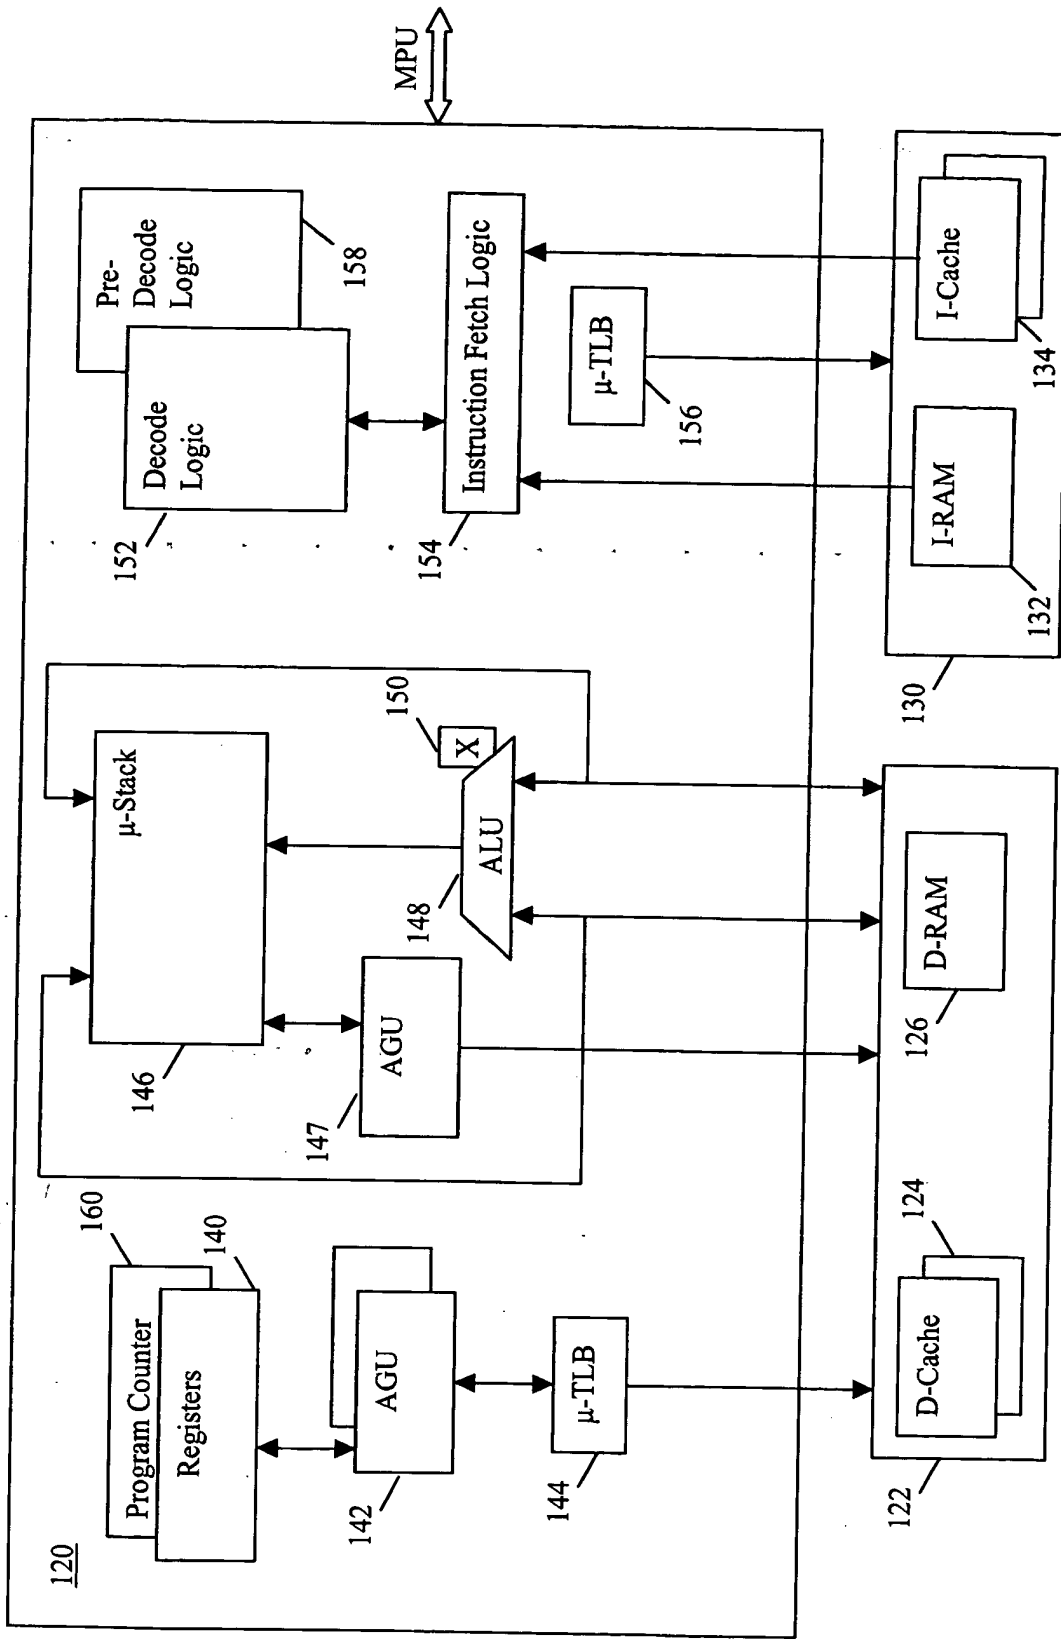


Figure 2

R0	General Purpose (GP)
R1	General Purpose (GP)
R2	General Purpose (GP)
R3	General Purpose (GP)
R4	Program Counter (PC)
R5	General Purpose/Local Variable Pointer (LV)
R6	Stack Pointer (SP)
R7	Top of Stack (ToS)
R8	General Purpose/Address Index 0 (AI0)
R9	General Purpose/Address Index 1 (AI1)
R10	General Purpose (GP)
R11	General Purpose (GP)
R12	Micro-program counter (micro-PC)
R13	General Purpose (GP)
R14	General Purpose/Indirect Register Index (IRI)
R15	Status and Control (ST)

Figure 3




Figure 4




Figure 5

